Minimax regret with imperfect *ex-post* knowledge of the state

Giacomo Bonanno

Department of Economics, University of California, Davis, USA
gfbonanno@ucdavis.edu

Abstract

We consider decision problems under complete ignorance and extend the minimax regret principle to situations where, after taking an action, the decision maker does not necessarily learn the state of the world. For example, if the decision maker only learns what the outcome is, then all she knows is that the actual state is one of the possibly several states that yield the observed outcome under the chosen action. We refer to this situation as *imperfect ex-post information*. We also extend the framework to encompass the possibility of less than the extreme pessimism that characterizes the minimax regret criterion.

Keywords: regret, minimax, pessimism, Hurwicz index

October 13, 2021

1 Introduction

A decision problem is typically framed in terms of three entities: (1) the available *actions* (also called acts), (2) the external facts over which the decision maker (henceforth, DM) has no control, usually referred to as *states*, and (3) the possible *outcomes* or consequences, over which the DM has well-defined preferences. An action is construed as a list of outcomes, one for each state.

The decision problem can then be represented as a matrix, as shown in Table 1 where the rows are labeled with actions and the columns with states and the entries in the matrix are the possible outcomes: in the example of Table 1 the outcomes are taken to be sums of money (changes in the DM's wealth).

		state		
		s_1	s_2	s_3
	a_1	\$10	\$10	\$2
action	a_2	\$1	\$30	\$1
	a_3	\$4	\$18	\$7

Table 1: A decision problem with monetary outcomes.

We will focus on decision criteria that are based on the notion of *regret aversion*, which rests on two assumptions.

The first assumption is that, once an action has been chosen and its consequences have been learned, the DM feels regret if an alternative action would have led to a better outcome. For example, if – in the decision problem represented in Table 1 – the DM chooses action a_1 and later learns that the state is s_2 then she experiences regret (assuming, of course, that she prefers more money to less), because she could have obtained an additional \$20 if she had chosen action a_2 instead of a_1. Thus, *ex-post* the DM experiences regret if she *knows* that she could have done better. Most of the literature on regret implicitly assumes that, after taking an action, the DM learns what the state is. This paper deals with the case where the *ex-post* knowledge acquired by the DM might be imperfect. For example, the DM might simply learn what the outcome is and not necessarily what the state is. If the observed outcome after taking action a is associated with several states, then all the DM learns is that the actual state is one of those states. For example, in the decision problem represented in Table 1, if the DM chooses action a_1 and receives \$10 then she learns that the state is either s_1 or s_2 and thus it is not the case that she knows that she could have done better with a different action: as far as she knows, the state *could be* s_1, in which case action a_1 is in fact the optimal action. On the other hand, if the DM takes action a_2 and receives \$1 then she learns that the state is either s_1 or s_3 and she knows that she would have done better with either a_1 or a_3. There is a sizeable empirical literature that discusses this issue, referred to as the "feedback" issue.

This literature is discussed in Section 2.

The second assumption on which the notion of regret aversion rests is that *ex ante* (that is, at the time of making a decision) the DM anticipates the *ex post* feelings of regret and chooses that action that best protects her from regret. While the first assumption has to do with the presence or absence of regret, the second assumption requires *measuring* the amount of regret experienced. This is done by postulating a von Neumann-Morgenstern utility function over the set of outcomes and constructing a "regret matrix" where the entry in the row corresponding to action a and state s is given as the difference between the largest utility that could have been achieved in state s (by possibly taking a different action) and the utility actually achieved. For example, in the situation represented in Table 1, if we assume that the DM's von Neumann-Morgenstern utility function is the identity function: $U(\$x) = x$, then the regret matrix is as shown in Table 2.

		state		
		s_1	s_2	s_3
	a_1	0	20	5
action	a_2	9	0	6
	a_3	6	12	0

Table 2: The regret matrix corresponding to Table 1 when the utility function is $U(\$x) = x$.

The most commonly used decision rule is *minimax regret*, introduced by Savage (1951) and later axiomatized by Milnor (1954), Puppe and Schlag (2009), Stoye (2011). The rule is based on the assumption that the DM is unable to assign probabilities to the states, a situation usually referred to as *complete ignorance*.[1] According to the minimax regret principle, one first determines, for every action, the largest possible regret and then chooses that action that minimizes the largest regret. For example, in Table 2 the largest regrets are: 20 for action

[1] On the other hand, regret theory – introduced by Bell (1982), Loomes and Sugden (1982; 1987) and axiomatized by Sugden (1993), Hayashi (2008), Diecidue and Somasundaram (2017) – takes as given a probability measure over the set of states and constructs a binary relation on actions based on expected regret. For general discussions of regret theory see Sugden (1985), Bleichrodt and Wakker (2015).

a_1, 9 for a_2 and 12 for a_3 and thus the minimax regret criterion picks action a_2.[2] Of course, this criterion is based on the implicit assumption that *ex post* the DM learns what the state is. The objective of this paper is to discuss how to adapt the minimax regret rule to situations of imperfect *ex-post* information about the state. This is done in Section 3, after reviewing – in Section 2 – the empirical literature that studies the effect of imperfect information on *ex ante* decisions.

An assumption underlying the minimax regret criterion is that, *ex ante*, the DM is *extremely pessimistic* about her potential ex-post regret: this is why she focuses on the largest possible regret associated with every action. After generalizing the standard minimax regret principle to the case of imperfect ex-post information, in Section 4 we further generalize the framework by introducing a *Hurwicz index of pessimism* (Hurwicz (1951)), which allows for less than extreme pessimism. Section 5 contains further discussion and a conclusion.

2 The role of feedback

As remarked above, a critical issue in modeling regret is the extent to which the DM, after her choice, is informed about the outcomes that would have resulted had she chosen differently. This issue has been explored in the experimental literature on "feedback-conditional" regret (Larrick and Boles (1995), Ritov (1996), Zeelemberg et al. (1996), Zeelemberg and Beattie (1997), Zeelemberg (1999)). The general finding in this literature is that people tend to prefer options which screen them from discovering the outcome of counterfactual choices: the anticipated negative emotion associated with regret is reduced or eliminated if people do not know the outcome of the forgone choice.

Larrick and Boles (1995) conducted an experiment in which a job recruit negotiated with a recruiter over the size of a sign-up bonus. The recruits were told that they had a single "best alternative to a negotiation agreement" (BATNA), for which they knew only the probability distribution over possible values. Half the subjects were told that they would find out the specific value of their BATNA only if they failed to reach an agreement in their current negotiation, while the other half were told that they would discover the value of their BATNA regardless of whether they reached an agreement. The authors found that the first feedback scenario led recruits to be more likely to accept a low offer: by being less ambitious, they increased the likelihood of reaching an

[2] Acker (1997) proposes a modification of the minimax regret rule based on taking into account, not only the maximum, but also the minimum foregone utility associated with each action and each state.

agreement and thereby (because of the lack of feedback on the actual value of the BATNA) shielded themselves from the possibility of experiencing regret. On the other hand, in the second feedback scenario recruits tended to hold out for a better offer, in order to avoid the possibility of learning later that they had settled for too little (relative to the BATNA).[3]

Ritov (1996) ran an experiment where subjects were presented with a choice between two independent monetary lotteries and asked which one they would choose under different scenarios.[4] The two lotteries had the same expected value, but one of the two involved higher risk and higher gains. In one scenario the subjects were told that only the chosen lottery would be resolved, while in the other scenario both the chosen lottery and the foregone lottery would be resolved and the outcome revealed. The author found that the subjects in the second scenario (complete resolution) chose the high-risk, high-gain option more often than the subjects who were told that only their selected lottery would be resolved. Thus expectations concerning the extent of uncertainty resolution played a significant role in *ex-ante* choices.

In a series of experiments Zeelemberg et al. (1996) demonstrated that the anticipation of regret, caused by the manipulation of expected feedback on foregone options, can promote not only risk-averse but also risk-seeking choices. In all the experiments participants were given a choice between a risky and a safe gamble. Participants always expected to learn the outcome of the chosen option, but in addition they could sometimes receive feedback on the foregone outcome. Participants who expected to receive feedback on the safe option, regardless of their choice, were likely to choose this option. Participants who expected to receive feedback on the risky option tended to choose the risky option. This pattern of choice was found in both high and low variance gambles, but was more pronounced in the latter.[5]

The experimental literature reviewed above shares two features: (1) subjects were only asked to choose between two alternatives and (2) the alternatives were monetary lotteries with specified probabilities, so that the decision problems were not characterized by complete ignorance. Furthermore, no general

[3] Although we have focused on the choices of the recruits, Larrick and Boles (1995) showed that the same behavior was exhibited by the recruiters.

[4] The subjects' choices, however, were only hypothetical since they were not actually paid the amounts of money associated with their chosen lotteries.

[5] Whereas most experiments focused on individual decision making, Zeelemberg and Beattie (1997) also considered decisions made in an interpersonal context, namely the ultimatum game. An overview of some of the empirical literature on the role of feedback in decision making is provided in Zeelemberg (1999).

theoretical framework was offered for modeling partial *ex-post* information (the only theoretical contribution I am aware of is Humphrey (2004), which will be discussed in Section 5). In the following section we provide such a framework.

3 A general model of partial *ex-post* information

What information the DM receives *ex post* should be part of the description of the decision problem. This can be done by associating with every action a a partition Π_a of the set of states Σ. For every state $s \in \Sigma$, we denote by $\Pi_a(s)$ the cell of Π_a that contains state s (thus $s \in \Pi_a(s)$ and, furthermore, if $s' \in \Pi_a(s)$ then $\Pi_a(s') = \Pi_a(s)$). The interpretation of $\Pi_a(s)$ is that, if the DM takes action a and the actual state is s, then all the DM learns is that the actual state is one of the states in $\Pi_a(s)$. A natural example of this is the case where, after taking action a, the DM *ex post* only learns what the outcome is and thus cannot distinguish between any two states that yield that outcome. For example, in the decision problem shown in Table 3 (which reproduces Table 1 with the addition of a column labeled "feedback"), if the DM only learns how much money she receives, then the partition associated with each action is shown in the last column.

		state			
		s_1	s_2	s_3	feedback
	a_1	\$10	\$10	\$2	$\{\{s_1, s_2\}, \{s_3\}\}$
action	a_2	\$1	\$30	\$1	$\{\{s_1, s_3\}, \{s_2\}\}$
	a_3	\$4	\$18	\$7	$\{\{s_1\}, \{s_2\}, \{s_3\}\}$

Table 3: The decision problem of Table 1 when the DM learns only the outcome of her choice and not necessarily the state.

There are many situations where one only learns the outcome of the choice made, but not what would have been the outcome of an alternative choice: a researched who submits a paper to a journal, and is notified of its acceptance, does not know whether the paper would have been accepted or rejected if she had submitted it to a more prestigious journal; a consumer who purchases an experience good is able to determine, through repeated use, how satisfied she is with it, but learns nothing about how it compares to a competitive good; an employer who hires a new recruit learns how competent and reliable the new hire is, but does not know whether an alternative applicant would have been better or worse, etc.

The case where *ex post* the DM always learns the state is a special case of this general framework: it corresponds to the situation where the partition associated with each action is the same, namely the finest partition (whose elements are singleton sets); we shall call this the case of *perfect information*. The expression *imperfect information* will be used to refer to the remaining cases where, for at least some action a, there is at least one element of Π_a that is not a singleton.

Consider the decision problem of Table 3 (with the feedback postulated in the last column) and suppose that the DM selects action a_1 and receives \$10. Will she experience regret? She does not know if the state is s_1, in which case action a_1 was indeed the optimal action, or s_2, in which case she would have received more money with an alternative action. Hence she considers it possible that she could have done worse and she also considers it possible that she could have done better. The important point is that she *does not know* that she could have done better and thus – we postulate – cannot experience regret. We state this as an assumption. First a definition. For every action a and state s we denote by $a(s) \in X$ the outcome that a yields in state s (X is the set of outcomes); furthermore, if $x, y \in X$ we write $x \succsim y$ to denote that the DM considers outcome x to be at least as good as outcome y and $x \succ y$ to denote that the DM considers x to be better than y (as is standard, we assume that the DM has a complete and transitive preference relation over the set of outcomes).[6]

Definition 3.1. Let Σ be the set of states, $S \subseteq \Sigma$ a set of states, and a and b two actions. We say that *b dominates a relative to S* if (1) $b(s) \succsim a(s)$, for all $s \in S$, and (2) for at least one $s' \in S$, $b(s') \succ a(s')$. We denote by $D(a, S) \subseteq A$ the (possibly empty) set of actions that dominate a relative to S (A is the set of actions).

Assumption 1. After taking action a, in order for the DM to experience regret at state s it must be the case that she knows that she would have done better with another action, that is, there must be an action b that dominates a relative to $\Pi_a(s)$ (recall that $\Pi_a(s)$ is the cell of the partition Π_a that contains state s). Thus, at state s, the DM regrets choosing action a if and only if $D(a, \Pi_a(s)) \neq \varnothing$.

[6] A binary relation $\succsim \subseteq X \times X$ is *complete* (or total) if, for all $x, y \in X$, either $x \succsim y$ or $y \succsim x$ (or both); it is *transitive* if for all $x, y, z \in X$, if $x \succsim y$ and $y \succsim z$ then $x \succsim z$. Strict preference, denoted by $x \succ y$, is defined as $x \succsim y$ and $y \not\succsim x$.

In the example of Table 3, the DM does not experience regret at state s_1 if she takes action a_1 (indeed, $D(a_1, \{s_1, s_2\}) = \emptyset$), whereas she does experience regret at state s_1 if she takes action a_2, since both a_1 and a_3 dominate a_2 relative to $\Pi_{a_2}(s_1) = \{s_1, s_3\}$ ($D(a_2, \{s_1, s_3\}) = \{a_1, a_3\}$).

The issue of when the DM experiences regret is separate from the issue of how regret should be measured when it is in fact experienced. For this we take the standard approach and postulate a von Neumann-Morgenstern utility function on the set of outcomes, $U : X \to \mathbb{R}$. Before dealing with the general case, let us consider the example of Table 3 and suppose that the DM has chosen action a_2 and receives \$1, thus inferring that the actual state is either s_1 or s_3; hence she comes to know that she would have done better with either a_1 or a_3. *Which of the two alternative actions does she wish she had taken?* Assume that the DM's utility function is the identity function: $U(\$x) = x$. If she had taken action a_1 she could have had an extra utility of $10 - 1 = 9$, if the state is s_1, or an extra utility of $2 - 1 = 1$, if the state is s_3; if pessimistic she will assess her regret of not taking action a_1 at 9, while if optimistic she will assess her regret of not taking action a_1 at 1. In keeping with the implicit assumption of extreme pessimism, which underlies the minimax regret principle, we will assume in this section that the DM focuses on the largest regret associated with not taking action a_1, namely 9; in Section 4 we will consider the possibility of less than extreme pessimism. Similarly, if she had taken action a_3 she could have had an extra utility of 3, if the state is s_1, or an extra utility of 6, if the state is s_3 and we assume that the DM assesses her regret of not taking action a_3 at 6, the larger of the two values. Finally, we take the larger of the two regret values, namely 9, as a measure of the regret experienced *ex post* in the state of knowledge $\{s_1, s_3\}$ after taking action a_2.

Thus, for every action a, we can associate a regret value to each cell of Π_a. For the decision problem of Table 3, the regret values are as shown in Table 4.

Finally, according to the minimax regret rule, we pick the action that has the lowest of the largest regrets. Thus in the case of Table 4 the chosen action is a_1 (note that, in the case of perfect information, the action chosen by the minimax regret rule is a_2: see Table 2).

	$\{s_1\}$	$\{s_2\}$	$\{s_3\}$	$\{s_1, s_2\}$	$\{s_1, s_3\}$
a_1	--	--	5	0	--
a_2	--	0	--	--	9
a_3	6	12	0	--	--

Table 4: The *ex-post* regret values for Table 3 when the utility function is $U(\$x) = x$.

Turning to the general framework, consider the case where the DM has chosen action a and the actual state is s. Two cases are possible:

1. $D(a, \Pi_a(s)) = \emptyset$, that is, there is no action that dominates a relative to $\Pi_a(s)$. Then it is not the case that the DM knows that she could have done better with some other action and thus we take regret to be 0.

2. $D(a, \Pi_a(s)) = \{b_1, \ldots, b_r\}$ with $r \geq 1$. Then the DM knows that she could have done better with any of the actions b_i $(i = 1, \ldots, r)$. For every $i = 1, \ldots, r$ let $R(b_i|a) = \max_{s' \in \Pi_a(s)} \{U(b_i(s')) - U(a(s'))\}$ be the regret of not taking action b_i (instead of a). Then we define the regret associated with taking action a at state s as $R(a, \Pi_a(s)) = \max_{i \in \{1,\ldots,r\}} \{R(b_i|a)\}$. [7] , [8]

Finally, applying the minimax regret principle, we select that action that minimizes the maximum regret. More precisely, denote by $\Gamma(a)$ the *ex-post* regret associated with action a and define it as follows: $\Gamma(a) = \max_{S \in \Pi_a} R(a, S)$. Let $\hat{A} \subseteq A$ be the set of actions selected by the minimax regret principle; then $a \in \hat{A}$ if and only if $\Gamma(a) \leq \Gamma(a')$ for all $a' \in A$.

[7]Note that if $s' \in \Pi_a(s)$ then $\Pi_a(s') = \Pi_a(s)$ and thus $R(a, \Pi_a(s')) = R(a, \Pi_a(s))$. Thus we could alternatively write $R(a, s)$ instead of $R(a, \Pi_a(s))$.

[8]Note also that if $\Pi_a(s) = \{s\}$, that is, if *ex post* the DM knows that the state is s, then, by Definition 3.1, $b \in D(a, \{s\})$ if and only if $U(b(s)) > U(a(s))$ and thus Point 2 above yields the standard measure of regret for the case of perfect information.

4 Allowing for less than extreme pessimism

As remarked above, the minimax regret principle reflects extreme *ex-ante* pessimism on the part of the DM, in that she assesses the *ex-post* regret of taking an action at its largest possible value. In this section we allow for less than extreme pessimism by assuming that the DM is characterized by a *Hurwicz index of pessimism* α (with $0 < \alpha \leq 1$) (Hurwicz (1951)) so that, both *ex post* and *ex ante*, she assesses regret as a convex combination of the largest and smallest regret values, namely α times the largest value plus $(1-\alpha)$ times the smallest value. The standard minimax regret criterion thus corresponds to the case where $\alpha = 1$ (together with perfect information).

In the case of perfect *ex-post* information, whenever – as is the case in the example of Table 1 – for every action there is a state where that action is optimal, the value of α makes no difference in terms of the selected action. Indeed, in such a case the assessed regret value for action a will be $\alpha m + (1-\alpha)0 = \alpha m$, where m is the maximum regret associated with a; hence minimizing αm leads to the same result as minimizing m. On the other hand, even in the case under consideration, the value of the index α does have an impact on the selected action if there is imperfect *ex-post* information. For example, in the case of Table 3, if the DM takes action a_2 then she will assess the regret of not taking action a_1 as $9\alpha + 1(1-\alpha) = 1 + 8\alpha$ and the regret of not taking action a_3 as $6\alpha + 3(1-\alpha) = 3 + 3\alpha$; thus, *ex post* she will asses her regret for having chosen action a_2 as $\max\{1+8\alpha, 3+3\alpha\} = \begin{cases} 3+3\alpha & \text{if } \alpha \leq \frac{2}{5} \\ 1+8\alpha & \text{if } \alpha > \frac{2}{5} \end{cases}$, so that the action that she will select *ex ante* is $\begin{cases} a_2 & \text{if } \alpha < \frac{1}{2} \\ a_1 \text{ or } a_2 & \text{if } \alpha = \frac{1}{2} \\ a_1 & \text{if } \alpha > \frac{1}{2} \end{cases}$.

In general, consider the case where the DM has chosen action a and the actual state is s.

1. *Ex-post* regret. Let $\rho(a, \Pi_a(s))$ be the DM's *ex-post* regret of taking action a knowing that the actual state belongs to the set $\Pi_a(s)$. Define $\rho(a, \Pi_a(s))$ as follows.

    (a) If $D(a, \Pi_a(s)) = \varnothing$ (that is, there is no action that dominates a relative to $\Pi_a(s)$) then $\rho(a, \Pi_a(s)) = 0$.

    (b) Suppose that $D(a, \Pi_a(s)) = \{b_1, \ldots, b_r\}$ with $r \geq 1$ (that is, the DM knows that she could have done better with any of the actions

b_i, $i = 1, \ldots, r$). Let $\alpha \in (0, 1]$ the the DM's index of pessimism. For every $i = 1, \ldots, r$ let $R(b_i|a) = \alpha \max_{s' \in \Pi_a(s)} \{U(b_i(s')) - U(a(s'))\} + (1 - \alpha) \min_{s' \in \Pi_a(s)} \{U(b_i(s')) - U(a(s'))\}$ be the regret of not taking action b_i (instead of a). Then $\rho(a, \Pi_a(s)) = \max_{i \in \{1,\ldots,r\}} \{R(b_i|a)\}$.

2. *Ex-ante* (or anticipated) regret. For every action $a \in A$, let $\Gamma(a)$ be the assessed *ex-ante* regret associated with action a, defined as follows: $\Gamma(a) = \alpha \max_{S \in \Pi_a} \{\rho(a, S)\} + (1 - \alpha) \min_{S \in \Pi_a} \{\rho(a, S)\}$.

Finally, let $\hat{A} \subseteq A$ be the set of actions selected by the generalized minimax regret principle; then $a \in \hat{A}$ if and only if $\Gamma(a) \le \Gamma(a')$, for every $a' \in A$. For example, consider the decision problem shown in Table 5 where the numbers are utilities, that is, the entry in the row labeled a_i and column labeled s_j is $U\left(a_i(s_j)\right)$.

		state			
		s_1	s_2	s_3	s_4
	a_1	1	1	2	2
action	a_2	2	1	3	1
	a_3	0	3	1	3
	a_4	1	4	1	4

Table 5: A decision problem where the entries are utilities.

Let us first consider the case of perfect information. Table 6 shows the values $\rho(a, \{s\})$.[9]

[9] These are the standard *ex-post* regret values, that is, $\rho(a, \{s\}) = \max_{a' \in A} \{U(a'(s))\} - U(a(s))$.

	s_1	s_2	s_3	s_4	*ex-ante* regret Γ
a_1	1	3	1	2	$\Gamma(a_1) = 3\alpha + (1-\alpha) = 1 + 2\alpha$
a_2	0	3	0	3	$\Gamma(a_2) = 3\alpha$
a_3	2	1	2	1	$\Gamma(a_3) = 2\alpha + (1-\alpha) = 1 + \alpha$
a_4	1	0	2	0	$\Gamma(a_4) = 2\alpha$

Table 6: The *ex-post* regret values $\rho(a, \{s\})$ for the decision problem of Table 5 in the case of perfect information. The corresponding *ex-ante* regret values $\Gamma(a_i)$ are shown in the last column.

Thus, in the case of perfect information, the set of actions selected by the generalized minimax regret criterion is as follows:

$$\hat{A} = \begin{cases} \{a_4\} & \text{if } \alpha < 1 \\ \{a_3, a_4\} & \text{if } \alpha = 1 \end{cases} .$$

Next we consider the imperfect-information case shown in Table 7 where the utilities are the same as in Table 5 and the feedback is as shown in the last column.

		state				
		s_1	s_2	s_3	s_4	feedback
	a_1	1	1	2	2	$\{\{s_1, s_2\}, \{s_3, s_4\}\}$
action	a_2	2	1	3	1	$\{\{s_1\}, \{s_2, s_4\}, \{s_3\}\}$
	a_3	0	3	1	3	$\{\{s_1\}, \{s_2, s_4\}, \{s_3\}\}$
	a_4	1	4	1	4	$\{\{s_1, s_3\}, \{s_2, s_4\}\}$

Table 7: The decision problem of Table 5 with the imperfect-information feedback shown in the last column.

Table 8 shows the *ex-post* regret values $\rho\,(a, \Pi_a(s))$ for the decision problem of Table 7.[10]

	$\{s_1\}$	$\{s_3\}$	$\{s_1, s_2\}$	$\{s_1, s_3\}$	$\{s_2, s_4\}$	$\{s_3, s_4\}$	*ex-ante* regret Γ
a_1	--	--	3α	--	--	0	$\Gamma(a_1) = 3\alpha^2$
a_2	0	0	--	--	3	--	$\Gamma(a_2) = 3\alpha$
a_3	2	2	--	--	1	--	$\Gamma(a_3) = 1 + \alpha$
a_4	--	--	--	$1 + \alpha$	0	--	$\Gamma(a_4) = (1 + \alpha)\alpha$

Table 8: The *ex-post* regret values for Table 7 and the corresponding *ex-ante* regret values $\Gamma(a_i)$ in the last column.

Thus the set of actions selected by the generalized minimax regret criterion is as follows:

$$\hat{A} = \begin{cases} \{a_1\} & \text{if } 0 < \alpha < \frac{1}{2} \\ \{a_1, a_4\} & \text{if } \alpha = \frac{1}{2} \\ \{a_4\} & \text{if } \frac{1}{2} < \alpha < 1 \\ \{a_3, a_4\} & \text{if } \alpha = 1 \end{cases}$$

5 Discussion and conclusion

In Section 2 we reviewed the empirical literature that highlighted the importance, for decision making, of the feedback about the state that the DM acquires *ex post*. To the best of my knowledge, the only attempt to model the feedback effect theoretically was made by Humphrey (2004). His approach, however, is substantially different from ours. The author adopts the framework of Loomes and Sugden (1987) where the DM is able to assign a probability p_j to each state s_j ($j = 1, \ldots, n$). The outcome of taking action a_i when the state is s_j is denoted by x_{ij}. Loomes and Sugden (1987) postulate that the utility of taking action a_i in state s_j depends not only on the actual outcome x_{ij} but also on the foregone outcome x_{kj} (from the foregone action a_k, $k \neq i$) and denote this utility by $M(x_{ij}, x_{kj})$; they then use the probabilities of the states to define

[10] For example, $D(a_1, \{s_1, s_2\}) = \{a_2, a_4\}$, $R(a_2|a_1) = \alpha$ and $R(a_4|a_1) = 3\alpha$ so that $\rho(a_1, \{s_1, s_2\}) = 3\alpha$. Similarly, $D(a_4, \{s_1, s_3\}) = \{a_1, a_2\}$, $R(a_1|a_4) = \alpha$ and $R(a_2|a_4) = 2\alpha + (1 - \alpha) = 1 + \alpha$ so that $\rho(a_4, \{s_1, s_3\}) = 1 + \alpha$.

a binary preference relation on the set of actions as follows: $a_i \succsim a_k$ if and only if $\sum_{j=1}^{n} M(x_{ij}, x_{kj})\, p_j \geq \sum_{j=1}^{n} M(x_{kj}, x_{ij})\, p_j$. Humphrey (2004) modifies this utility function in order to take into account the feedback that the DM receives about the state. The author interprets Loomes and Sugden (1987)'s utility function $M(x_{ij}, x_{kj})$ (which he denotes by $m(x_{ij}, x_{kj})$) as the

> "anticipated utility of having x_{ij} and missing out on x_{kj} under a particular state of the world where the outcome of the chosen act x_{ij} is *fully revealing of the state of the world*. In this case receiving x_{ij} reveals that the outcome of the foregone act is x_{kj}." [(Humphrey 2004, p.845), emphasis in the original.]

To the function $M(x_{ij}, x_{kj})$ the author adds a second function $\mu(x_{ij}, x_{kj})$ which he interprets as the

> "modified anticipated utility of having x_{ij} and missing out on x_{kj}, but where having x_{ij} will *not be fully revealing of the state of the world*. In this case, the decision-maker has anticipated a state of the world under which they will receive x_{ij} and forego x_{kj}, but actually receiving x_{ij} does not reveal x_{kj} (as opposed to some other outcome, say, z) as the outcome of the foregone act." [*ibidem*, emphasis in the original.]

The author then imposes several restrictions on the relationship between these two functions and shows that some of the conclusions reached in the empirical literature can be reversed by an appropriate choice of the additional function $\mu(x_{ij}, x_{kj})$. However, the author does not offer a general theory of what it means for an outcome "not to be fully revealing of the state of the world"; in particular, if – after taking action a_i – outcome x_{ij} occurs (so that the state is s_j), what restrictions are there on the alternative states to be considered? Furthermore, he offers no explanation of how the function μ varies with alternative forms of feedback and how such a function should be constructed or elicited.

We considered the case of complete ignorance (where the DM is not able to assign probabilities to the states) and presented a general framework for adapting the minimax regret principle to situations where, after taking an action, the DM does not necessarily learn what the actual state is. For example, if – after taking action a – the DM *ex post* only learns that the outcome is x_0, then she will not be able to distinguish between any two states that yield outcome x_0 under action a: the set of states that she considers possible is

$S_0 = \{s \in \Sigma : a(s) = x_0\}$. We postulated that, in order to experience regret, the DM must *know* that she would have done better with a specific different action, that is, there must be an alternative action b that dominates a relative to the set of states S_0, in the sense that $b(s) \succsim a(s)$ (outcome $b(s)$ is at least as good as outcome $a(s)$), for every $s \in S_0$, and, furthermore, for at least one $s' \in S_0$, $b(s') \succ a(s')$ ($b(s')$ is better than $a(s')$). We take it that, for the DM, regretting choosing action a corresponds to a statement of the form *"in light of what I know now*, I wish I had chosen action b instead of a". If, according to her *ex-post* knowledge, the DM thinks that the state could be s_1 or could be s_2, then she cannot consistently think "I wish that I had chosen action b_1 if the state is s_1 and action b_2 if the state is s_2", since she cannot distinguish between s_1 and s_2; in other words, *the wished hypothetical alternative action cannot be a function of contingencies that the DM cannot tell apart ex post*.

Having identified a necessary and sufficient condition for the experience of *ex-post* regret, we then dealt with the issue of how to measure regret. This step involves assessing the desirability of every alternative action that qualifies (that is, that dominates the chosen action relative to the acquire *ex-post* knowledge), assigning a regret value to each such action and then picking the one with the largest regret value. In Section 4 we further generalized the framework by allowing for less than the extreme degree of pessimism that underlies the minimax regret principle: we assumed that the DM is characterized by a Hurwicz index of pessimism $\alpha \in (0, 1]$. The generalized minimax regret principle reduces to the standard minimax principle when (1) there is perfect information (that is, the DM *ex post* learns the state) and (2) $\alpha = 1$.

References

M. H. Acker. Tempered regrets under total ignorance. *Theory and Decision*, 42: 207–213, 1997.

D. E. Bell. Regret in decision making under uncertainty. *Operations Research*, 30:961–981, 1982.

H. Bleichrodt and P. P. Wakker. Regret theory: a bold alternative to the alternatives. *Economic Journal*, 125:493–532, 2015.

E. Diecidue and J. Somasundaram. Regret theory: A new foundation. *Journal of Economic Theory*, 172:88–119, 2017.

T. Hayashi. Regret aversion and opportunity dependence. *Journal of Economic Theory*, 139:242–268, 2008.

S. J. Humphrey. Feedback-conditional regret theory and testing regret-aversion in risky choice. *Journal of Economic Psychology*, 25:839–857, 2004.

L. Hurwicz. The generalized Bayes minimax principle: a criterion for decision making under uncertainty. Technical report, Cowles Commission Discussion Paper: Statistics No. 355, 1951. URL https://cowles.yale.edu/sites/default/files/files/pub/cdp/s-0355.pdf.

R. Larrick and T. Boles. Avoiding regret in decisions with feedback: A negotiation example. *Organizational Behavioral and Human Decision Processes*, 63: 87–97, 1995.

G. Loomes and R. Sugden. Regret theory: an alternative theory of rational choice under uncertainty. *Economic Journal*, 92:805–824, 1982.

G. Loomes and R. Sugden. Some implications of a more general form of regret theory. *Journal of Economic Theory*, 41:270–287, 1987.

J. Milnor. Games against nature. In R. Thrall, C. Coombs, and R. Davis, editors, *Decision Processes*, pages 49–59. Wiley, New York, 1954.

C. Puppe and K. H. Schlag. Choice under complete uncertainty when outcome spaces are state dependent. *Theory and Decision*, 66:1–16, 2009.

I. Ritov. Probability of regret: Anticipation of uncertainty resolution in choice. *Organizational Behavioral and Human Decision Processes*, 66:228–236, 1996.

L. J. Savage. The theory of statistical decision. *Journal of the American Statistical Association*, 46:55–67, 1951.

J. Stoye. Axioms for minimax regret choice correspondences. *Journal of Economic Theory*, 146:2226–2251, 2011.

R. Sugden. Regret, recrimination and rationality. *Theory and Decision*, 19:77–99, 1985.

R. Sugden. An axiomatic foundation for regret theory. *Journal of Economic Theory*, 60:159–180, 1993.

M. Zeelemberg. Anticipated regret, expected feedback and behavioral decision making. *Journal of Behavioral Decision Making*, 12:93–106, 1999.

M. Zeelemberg and J. Beattie. Consequences of regret aversion 2: Additional evidence for effects of feedback on decision making. *Organizational Behavioral and Human Decision Processes*, 72:63–78, 1997.

M. Zeelemberg, J. Beattie, J. van der Pligt, and N. de Vries. Consequences of regret aversion: Effects of feedback on risky decision making. *Organizational Behavioral and Human Decision Processes*, 65:148–158, 1996.